

REMARKS

Claims 1-71 are pending in the application, of which Claims 1, 24, 25, 28, 32, 34, 57, 58, 61, 65, and 67-71 are independent. All claims have been rejected. These rejections are respectfully traversed and reconsideration is requested. All claims are believed to be in condition for allowance.

Rejections Under 35 U.S.C. 102

Claims 1-5, 8, 9, 11, 14, 15, 23, 25, 28, 32-38, 41, 42, 44, 47, 48, 56, 58, 61, 65-67 and 69-71 have been rejected under 35 U.S.C. 102(e) as being anticipated by Li et al. (U.S. Patent No. 6,560,230 B1). This rejection is respectfully traversed and reconsideration is requested.

The present application is directed to a scheduler for scheduling the order in which packets leave a network router. The network router can have several queues storing data packets to be transmitted. The scheduler, which selects a queue from which a packet is forwarded, holds scheduling values associated with the queues. The scheduling values may include schedule transmission rates according to a constant bit rate (CBR) service guarantee or, for example, represent theoretical transmission times using a weighted-fair-queuing (WFQ) scheduling policy. Scheduling values from one or more scheduling methods are compared in a selection network to select the packets to be forwarded.

In accordance with one aspect of the invention, the selection network can be a tree structure where each leaf of the tree structure represents a scheduling value of a queue. Internal nodes of the tree structure represent winners in comparisons of scheduling values of sibling nodes of the tree structure. Alternatively, the selection network may be a sorting network, as illustrated in FIG. 11, by which the scheduling values are compared to order to the queues by scheduling priority.

In another aspect of the invention, variable-length packets are present in the system and the scheduler takes into account the lengths of the packets when scheduling them for transmission (see FIGs. 12 and 13).

In yet another aspect of the invention, the packets selected for forwarding are stored in intermediate queues for yet another selection process (see FIG. 17). The intermediate queues

may, for example, store packets sorted according to different scheduling methods (Specification, page 25, lines 20-29).

In yet another aspect of the invention, indicators are associated with the queues. Such indicators may be used to disable the queues so they do not participate in scheduling. (See FIG. 7).

Li *et al.* teaches a packet scheduling policy that involves allocating bandwidth among classes of packets where the classes are arranged as a tree structure (FIG. 4). A scheduling engine is associated with each node of this policy tree. Leaf nodes of the tree select among the head packets of multiple input queues (FIG. 5A, column 11, lines 3-52). Each packet has an associated start time and finish time, and each leaf scheduling engine selects from among the eligible packets the packet with the earliest finish time. The selected packet is held by the leaf node until consumed by a higher-level node in the policy tree. After selection, the start and finish times of a packet are recomputed by the leaf node depending on the amount of bandwidth allocated to the class of packets represented by that leaf node (column 12, lines 4 to 14). The scheduling engines associated with internal nodes of the policy tree operate in a similar manner except they consider the packets held by lower nodes in the tree (rather than the packets at the heads of queues) and they take into account the priority of the child node holding the packet in selecting the eligible packet to send.]

Li *et al.* does not teach or suggest a selection network by which the scheduling values are compared to select packets to forward as is claimed, for example, in the independent Claims 1 and 67. While the system of Li *et al.* does compare scheduling values to select eligible packets to dequeue at leaf nodes (column 13, lines 26-47), it does not use a network to perform this function. Instead, separate scheduling engines are responsible for the task. The scheduling engines of Li *et al.* are not described in detail. All that is stated (column 14, lines 21-30) is that "such a scheduler could be implemented in as software...or it could be implemented as a hardware device."

Li *et al.* does disclose a policy tree (Fig. 4), however, the policy tree does not compare scheduling values associated with queues and is not used thus. Instead, it meters bandwidth according to a policy. Only the leaf nodes of the policy tree consider the scheduling values

associated with the input queues, while internal nodes compare values according to the class represented by that node.

Nowhere does Li *et al.* teach or suggest using the selection network to determine which packets to forward next. Therefore, Claims 1 and 67 are not anticipated by Li *et al.*, and the rejection should be withdrawn. Dependent Claims 2-23 depend on Claim 1 and are not anticipated by Li *et al.* for at least the same reasons as above.

The policy tree of Li *et al.* does not anticipate the tree-structured selection network of Claim 2 because internal nodes of the policy tree do not compare scheduling values associated with the queues. Therefore, dependent Claims 2-22 are patentable in view of Li *et al.*

Dependent Claim 3 recites limiting scheduler comparisons to the changed scheduling values. Li *et al.* does not teach or suggest a mechanism for limiting comparisons to scheduling values that have changed. The passage referred to by the Examiner (column 12, lines 40-59) discusses assignment of priority to leaf classes of the policy tree. It does not mention limiting comparisons to values that have changed. Therefore, dependent Claim 3 is not anticipated by Li *et al.*

Dependent Claim 4 recites internal nodes of the tree structure that store scheduling values from winning sibling nodes. Li *et al.* does not teach or suggest storing scheduling values associated with the queues in the internal nodes. Instead, the system of Li *et al.* computes new start and finish times for each packet at each node of the tree according to the scheduling policy. Therefore, dependent Claim 4 is not anticipated by Li *et al.* Furthermore, Li *et al.* does not describe storing identities of the winning nodes in the internal nodes, and, for at least that reason, Claim 5 is not anticipated by Li *et al.*

Dependent Claim 8 recites a pipelined scheduler. Pipelining is not taught in Li *et al.* The passage cited by the Examiner (column 12, lines 40-59) does not describe pipelining, but rather deals with the use of priority in scheduling. For at least this reason, Claim 8 and claims 9 and 10, which depend on it, are patentable in view of Li *et al.*

Dependent Claim 11 recites an identification of the path to a winning leaf at internal nodes of the tree. That feature is not taught or suggested by Li *et al.* The passage cited by the Examiner (column 9, lines 7-34) refers to the function of a leaf-node scheduler and not to the

identification of the path to the winning leaf. For at least this reason, dependent Claim 11 is patentable in view of Li *et al.*

Dependent Claim 23 further limits the scheduler to a sorting network. Nowhere does Li *et al.* describe or suggest using the sorting network to compare packet scheduling values. Therefore, dependent Claim 23 is patentable in view of Li *et al.*

Li *et al.* does not teach or suggest scheduling packets using two different scheduling methods, each with its own scheduling values, as is claimed in the independent Claims 25, 58 and 69. Instead, the system of Li *et al.* uses a single scheduling method based on start time, finish time, and priority for all packets. Using a second method of scheduling is not suggested in Li *et al.*, nor is any such method described. Therefore, independent Claims 25, 58 and 69 are not anticipated by Li *et al.* Dependent Claims 26-27, and 59-60 depend on the independent claims 25 and 58, respectively, and are not anticipated by Li *et al.* for at least the same reasons as above.

Independent Claims 28 and 70 recite, among other limitations, scheduling variable-length packets, which are not taught or suggested in Li *et al.* All packets of Li *et al.* have a fixed length that is known at the time they enter the system. Claim 28, in comparison, recites updating the scheduling value of a queue based on a variable length of a packet in the queue. Therefore, independent Claims 28 and 70 are not anticipated by Li *et al.* Dependent claims 29-31 depend on the independent Claim 28, and are patentable in view of Li *et al.* for at least the same reasons as above.

Independent Claims 32, 65 and 71 recite, among other limitations, a set of intermediate queues. Li *et al.* does not teach or suggest intermediate queues; instead, each internal node of the system of Li *et al.* holds only a single packet, not a queue of packets. Therefore, Claims 32, 65, and 71 are not anticipated by Li *et al.* Dependent Claims 33-56 and 66 depend on the independent Claims 32 and 65, respectively, and are not anticipated by Li *et al.* for at least the same reasons as above.

Rejections Under 35 USC § 103

Claims 6-7, 10, 12-13, 16-22, 24, 26-27, 29-31, 39-40, 43, 45-46, 49-55, 57, 59-60, 62-64, and 68 have been rejected under 35 U.S.C. 103(a) as being unpatentable over Li *et al.* in view of Lauer *et al.* (US 5,455,825). This rejection is respectfully traversed and reconsideration is requested.

Lauer *et al.* teaches an ATM cell scheduler in which each arriving cell is assigned a tag (Abstract). A search circuit 70 (Figure 1) is then used to search for the lowest-valued tag associated with a particular output port. Various embodiments are described in which the search is performed in a linear manner (Figure 3), as a binary tree (Figures 4 and 5), and in a hybrid manner – as a tree of linear searches (Figure 6).

Li *et al.* and Lauer *et al.* describe scheduling systems with contrasting scheduling policies, and a person skilled at the art at the time of the invention would not have a motivation to combine the teachings of the references. Li *et al.* describes a policy-tree based scheduler where scheduling decisions are based on class, and on start and finish times that are computed at each node of a tree. Lauer *et al.*, in contrast, describes a tag-based scheduler where scheduling decisions are made based on a single tag value that is assigned to cells as they enter the system. Moreover, the system of Li *et al.* selects packets from multiple queues, while Lauer *et al.* uses a single shared memory. Therefore, the combination of the two references is improper and would not have occurred to one skilled in the art at the time of the invention.

Even if combined, the references would not suggest the present invention. As discussed above, dependent claims 6-7, 10, 12-13, 16-22, 26-27, 29-31, 39-40, 43, 45-46, 49-55, 59-0, and 62-64 are not anticipated by Li *et al.* Lauer *et al.* fails to supplement Li *et al.* in such a way as to render those claims obvious. In particular, Lauer *et al.* does not teach a scheduling network, , scheduling variable-length packets, or using more than one methodology for deriving scheduling values. Therefore, the above-listed claims are not anticipated by the combination of Li *et al.* and Lauer *et al.* and the rejection should be withdrawn.

Neither Li *et al.* nor Lauer *et al.* teach or suggest indicators associated with the queues to disable the queues as recited, for example, in the independent Claims 24, 57, and 68. Neither the references taken separately, nor the combination of the two, suggest disabling some queues and

scheduling packets only in the queues that are not disabled. *Li et al.* does not teach or suggest indicators associated with the queues and *Lauer et al.* does not teach multiple queues at all. Therefore, independent Claims 24, 57, and 68 are not obvious in view of the combination of *Li et al.* and *Lauer et al.* and the rejection should be withdrawn.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 

James M. Smith

Registration No. 28,043

Telephone: (978) 341-0036

Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: 